```
Set      Items    Description
S1     5086700    CLUSTER? OR GROUP? OR PARTITION? OR CATEGORI?
S2        2061    MATERIAL?()VIEW? OR (STORE? OR SAVE? OR CACHE?)(N)QUER? OR
                  (AUXILIAR? OR BASE)()RELATION?
S3         256    S1 AND S2
S4          21    S3 AND (NODE? OR LOCATION? OR PROCESSOR?)
S5          13    RD (unique items)
S6          12    S5 NOT PY>2001
File    8:Ei Compendex(R) 1970-2004/Apr W4
          (c) 2004 Elsevier Eng.  Info. Inc.
File   35:Dissertation Abs Online 1861-2004/Apr
          (c) 2004 ProQuest Info&Learning
File  202:Info. Sci. & Tech. Abs. 1966-2004/Feb 27
          (c) 2004 EBSCO Publishing
File   65:Inside Conferences 1993-2004/May W1
          (c) 2004 BLDSC all rts. reserv.
File    2:INSPEC 1969-2004/Apr W4
          (c) 2004 Institution of Electrical Engineers
File   94:JICST-EPlus 1985-2004/Apr W2
          (c)2004 Japan Science and Tech Corp(JST)
File  111:TGG Natl.Newspaper Index(SM) 1979-2004/May 07
          (c) 2004 The Gale Group
File  233:Internet & Personal Comp. Abs. 1981-2003/Sep
          (c) 2003 EBSCO Pub.
File    6:NTIS 1964-2004/May W2
          (c) 2004 NTIS, Intl Cpyrght All Rights Res
File  144:Pascal 1973-2004/Apr W4
          (c) 2004 INIST/CNRS
File  434:SciSearch(R) Cited Ref Sci 1974-1989/Dec
          (c) 1998 Inst for Sci Info
File   34:SciSearch(R) Cited Ref Sci 1990-2004/May W1
          (c) 2004 Inst for Sci Info
File   99:Wilson Appl. Sci & Tech Abs 1983-2004/Mar
          (c) 2004 The HW Wilson Co.
File   95:TEME-Technology & Management 1989-2004/Apr W3
          (c) 2004 FIZ TECHNIK
```

03917270   E.I. No: EIP94081361123
   **Title:  Scheduling the allocation  of data fragments in a distributed database environment: A machine learning approach**
   Author: Chaturvedi, Alok R.; Choubey, Ashok K.; Roan, Jinsheng
   Corporate Source: Purdue Univ, West Lafayette, IN, USA
   Source:  IEEE Transactions on Engineering Management v 41 n 2 May 1994. p 194-207
   Publication Year: 1994
   CODEN: IEEMA4    ISSN: 0018-9391
   Language: English
   Document Type: JA; (Journal Article)    Treatment: G; (General Review); T; (Theoretical)
   Journal Announcement: 9409W5
   Abstract: Different database fragmentation and allocation strategies have been proposed to partially replicate data in a **partitioned** , distributed database (DDB) environment. The replication strategies include database snapshots, **materialized   views** , and quasi-copies. These strategies are ?static' and do not adapt to the changes in the data usage patterns. Furthermore, they often require expensive update synchronizations to maintain data consistency and do not exploit the knowledge embedded in the query history. This paper describes a machine learning based time invariant fragmentation method (MLTIF) that acquires knowledge about the data usage patterns for each **node** . Based on this knowledge, MLTIF designs time invariant fragments (TIF) and schedules its allocation and selective update for a specified time period. Simulation is used to compare the effectiveness of the MLTIF approach with that of full replication, **materialized   views** , and non replication strategies. Initial results indicate that for most normal operating conditions, the MLTIF approach can be effective. (Author abstract) 21 Refs.
   Descriptors: *Storage allocation (computer); Distributed database systems ; Learning systems; Scheduling; Data processing; Large scale systems; Reliability
   Identifiers: Data fragment allocation; Machine learning based time invariant fragmentation method; Time invariant fragmentation
   Classification Codes:
   723.1  (Computer Programming); 723.2  (Data Processing); 723.4 (Artificial Intelligence); 922.2  (Mathematical Statistics)
   723  (Computer Software); 922  (Statistical Methods)
   72  (COMPUTERS & DATA PROCESSING); 92  (ENGINEERING MATHEMATICS

**QUERY PROCESSING IN SPATIAL DATABASE SYSTEMS: DECLUSTERING AND** CLUSTERING **TECHNIQUES**
    Author:  RAVADA, SIVAKUMAR
    Degree:  PH.D.
    Year:    1997
    Corporate Source/Institution:  UNIVERSITY OF MINNESOTA (0130)
    Source:  VOLUME 58/05-B OF DISSERTATION ABSTRACTS INTERNATIONAL.
             PAGE 2515.  94 PAGES
    Descriptors:  COMPUTER SCIENCE
    Descriptor Codes:  0984

     The research question in this thesis concerns how to parallelize the
spatial range and join query processing in order to support a high
performance spatial database application. Data **partitioning** for the range
query operation involves declustering of spatial data, while data
**partitioning** for the spatial join involves **clustering** of spatial data.
If the static **partitioning** methods fail to equally distribute the load
among different **processors** , the load-balance may be improved by
redistributing parts of the data to idle **processors** using Dynamic
Load-Balancing (DLB) techniques.
     In this thesis, we provide a framework for declustering collections of
extended spatial objects by identifying the following key issues: (i) the
work-load metric, (ii) the spatial-extent of the work-load, (iii) the
distribution of the work-load over the spatial-extent, and (iv) the
declustering method. We identify and experimentally evaluate alternatives
for each of these issues.
     In addition, we also provide a framework for dynamically balancing the
load between different **processors** . We experimentally evaluate the
proposed declustering and load-balancing methods on a distributed memory
MIMD machine (Cray T3D) and shared-memory machine (SGI Challenge).
Experimental results show that the spatial-extent and the work-load metric
are important issues in developing a declustering method. Experiments also
show that the replication of data is usually needed to facilitate dynamic
load-balancing, as the cost of local processing is often less than the cost
of data transfer for extended spatial objects. In addition, we also show
that the effectiveness of dynamic load-balancing techniques can be improved
by using declustering methods to determine the subsets of spatial objects
to be transferred during run-time.
     A spatial join is often performed in two steps: a filter step and a
refinement step. In this thesis, we focus on the refinement step of the
spatial join. The refinement step of the spatial join takes as input a
sequence of pairs of tuples and checks each tuple to see if the join
predicate is satisfied for that tuple. This is similar to the join index
processing done in traditional relational databases. We develop min-cut
graph **partitioning** based methods for join processing using a join index.
We use min-cut graph **partitioning** as a new heuristic for solving the page
access sequence problem for fixed size buffer in sequential systems. We
show that the number of page accesses needed to compute a join using join
index in a fixed buffer environment is bounded by the sum of sizes of the
**base relations** and the size of the cut-set of the page connectivity
graph. Since the min-cut graph **partitioning** aims to minimize the size of
the cut-set, this proposed heuristic is a direct method. Experiments with
benchmark data sets show that the graph-**partitioning** based heuristic
outperforms the existing methods, particularly when join selectivity is
high and buffer space is small. (Abstract shortened by UMI.)

01314895   ORDER NO: AAD93-23023
**DATABASE PROCESSING ON AN N-DIMENSIONAL MESH (SUPERCOMPUTING)**
   Author:   JHANG, HYOUNG
   Degree:   PH.D.
   Year:     1992
   Corporate Source/Institution:   UNIVERSITY OF KANSAS (0099)
   Source:   VOLUME 54/06-B OF DISSERTATION ABSTRACTS INTERNATIONAL.
         PAGE 3254.   288 PAGES
   Descriptors:   ENGINEERING, ELECTRONICS AND ELECTRICAL; COMPUTER SCIENCE
   Descriptor Codes:   0544; 0984

In this dissertation, an array **processor** connected in an n-dimensional mesh is evaluated for very large relational database applications. Each **node** is assumed to have substantial processing capability with substantial amount of memory. By **partitioning** each **base relation** among all **nodes**, and by storing them directly within primary memory of a memory-resident database machine, the bottleneck of slow secondary storage access is eliminated. Whenever inter- **node** operations are required, local **partitions** are routed to other **nodes**.

In this dissertation, an n-dimensional mesh is studied for SIMD database processing. Database algorithms are classified into five **categories** based on their routing requirements: (1) database algorithms which do not require any routing such as select; (2) database algorithms that require partial routing without order such as tuple balancing; (3) database algorithms that require partial routing with order such as project, elimination of duplicates, and sort; (4) database algorithms that require full routing such as nested-loop, sort-merge join algorithms, and cartesian product; (5) database algorithms that require random routing such as hash-based join and hash-based aggregate algorithms. Routing algorithms appropriate for these different classes of database algorithms are presented, and their performances are analyzed. Their performances on an n-dimensional mesh are compared with their performances on a binary cube, which is a subset of the n-dimensional mesh. Incremental expansion is studied on an n-dimensional mesh and its impact on performance is measured.

The design of a large parallel computer is often limited by the longest interconnection. In an n-dimensional mesh, the method of interconnection, individual interconnection length and **node** architecture remain identical as the computer is expanded. A custom **node** architecture maximizes processing and communication parallelism. Another **node** architecture uses standard microprocessor components.

Several n-dimensional meshes are connected through serial communication links in the (n + 1)th dimension to support even larger database applications. Each n-dimensional mesh may execute a different database operation or query resulting in an MIMD database computer. Two query execution strategies are presented and their performance are studied using simulation models.

6611426    INSPEC Abstract Number: C2000-07-7120-033
 Title: **Active replication and update of content in electronic commerce**
   Author(s): Chaturvedi, A.R.; Choubey, A.K.; Jinsheng Roan
   Author Affiliation: Krannert Graduate Sch. of Manage., Purdue Univ., West
Lafayette, IN, USA
   Journal: International Journal of Electronic Commerce    vol.4, no.3
p.45-67
   Publisher: M.E. Sharpe,
   Publication Date: Spring 2000  Country of Publication: USA
   CODEN: IJECFE  ISSN: 1086-4415
   SICI: 1086-4415(200021)4:3L.45:ARUC;1-X
   Material Identity Number: G303-2000-002
   U.S. Copyright Clearance Center Code: 1086-4415/2000/$9.50+0.00
   Language: English    Document Type: Journal Paper (JP)
   Treatment: Practical (P)
   Abstract:  This  paper  describes a strategy for actively replicating and
updating content  for electronic commerce. Active replication and updating
of content  is  achieved by intelligent agents (IA) using a time-invariant
fragmentation  approach  to  **partitioning**  and  replicating data  in  a
distributed computing environment. Taking into account the time-sensitivity
property  of date, IAs derive time-invariant fragments for their respective
 **nodes**   . From the query history. A time-invariant fragment (TIF) is that
portion  of the database whose values do not change during a specified time
interval. The algorithm that IAs use in creating TIFs for each **node** , for
a  given  time  interval,  is presented. The active replication approach is
compared with three other approaches, full-replication, nonreplication, and
 **materialized    view** ,  in  terms  of  data  transmission costs. Results
indicate  that  the  active  approach  can be most effective for electronic
commerce  because of the high percentage of modification queries, the large
size of the network, and the great number of transactions.  (27 Refs)
   Subfile: C
   Descriptors: concurrency control; electronic commerce; replicated
databases; software agents
   Identifiers: electronic commerce; replication; updating of content;
intelligent agents; distributed computing environment; replicating data;
time-invariant fragment; active replication
   Class Codes: C7120  (Financial computing); C6160B (Distributed databases)
; C6170  (Expert systems and other AI software and techniques)

5473273    INSPEC Abstract Number: C9702-4250-015
 **Title: On the optimality of degree of declustering**
   Author(s): Sheu, S.; Hua, K.A.; Cai, Y.
   Author Affiliation: Dept. of Comput. Sci., Univ. of Central Florida,
Orlando, FL, USA
   Conference Title: Database and Expert Systems Applications. 7th
International Conference, DEXA '96 Proceedings    p.865-74
   Editor(s): Wagner, R.R.; Thoma, H.
   Publisher: Springer-Verlag, Berlin, Germany
   Publication Date: 1996  Country of Publication: Germany    xv+921 pp.
   ISBN: 3 540 61656 X    Material Identity Number: XX96-03737
   Conference Title: Database and Expert Systems Applications. 7th
International Conference, DEXA '96 Proceedings
   Conference Date: 9-13 Sept. 1996    Conference Location: Zurich,
Switzerland
   Language: English    Document Type: Conference Paper (PA)
   Treatment: Theoretical (T)
   Abstract: Presents a data **partitioning** technique for shared-nothing
database systems. A unique feature of our scheme is that it organizes a
multicomputer system into **groups** of even numbers of processing **nodes** ;
and each relation is assigned to one of these **groups** in such a way to
minimize contention among concurrent queries. Thus, a fixed degree of
declustering is used for all **base    relations** in this scheme. Our
simulation results demonstrate that this approach provides significantly
better performance than conventional methods which independently determine
a degree of declustering for each of the **base    relations** . These schemes
totally ignore the requirement of inter-query parallelism. Obviously, an
appropriate degree of declustering represents a good trade-off between
inter-query and intra-query parallelism for our strategy. To investigate
this issue, we perform extensive simulations to study the effect of various
system and workload parameters on the optimality of the degree of
declustering. We found that it is influenced primarily by the parallel
processing overhead. With this finding, we develop a mathematical model to
determine the optimal degree of declustering for a given system. (10 Refs)
   Subfile: C
   Descriptors: database theory; distributed databases; optimisation;
parallel processing; query processing; resource allocation; simulation;
very large databases
   Identifiers: declustering degree optimality; data **partitioning**
technique; shared-nothing database systems; multicomputer system;
processing **nodes** ; contention minimization; concurrent queries; **base
relations** ; simulation; performance; inter-query parallelism; intra-query
parallelism; system parameters; workload parameters; parallel processing
overhead
   Class Codes: C4250  (Database theory); C6160B (Distributed databases);
C1180  (Optimisation techniques)

04244127   INSPEC Abstract Number: C9211-4250-009
 **Title: Parallel hierarchical evaluation of transitive closure queries**
  Author(s): Houtsma, M.A.W.; Cacace, F.; Ceri, S.
  Author Affiliation: Dept. of Comput. Sci., Twente Univ., Enschede, Netherlands
  Conference Title: Proceedings of the First International Conference on Parallel and Distributed Information Systems (Cat. No.91TH0393-4)    p. 130-7
  Publisher: IEEE Comput. Soc. Press, Los Alamitos, CA, USA
  Publication Date: 1991  Country of Publication: USA    xiii+292 pp.
  ISBN: 0 8186 2295 4
  U.S. Copyright Clearance Center Code: 0 8186 2295 4/91$01.00
  Conference Sponsor: IEEE; ACM; Florida Int. Univ.; Carnot Project
  Conference Date: 4-6 Dec. 1991    Conference Location: Miami Beach, FL, USA
  Language: English    Document Type: Conference Paper (PA)
  Treatment: Theoretical (T)
  Abstract: Presents a new approach to parallel computation of transitive closure queries using a semantic data fragmentation. Tuples of a large **base relation** denote edges in a graph, which models a transportation network. A fragmentation algorithm is proposed which produces a **partitioning** of the **base relation** into several fragments such that any fragment corresponds to a subgraph. One fragment, called high-speed fragment, collects all edges which guarantee maximum speed. Thus, the fragmentation algorithm induces a hierarchical relationship between the high-speed fragment and all other fragments. With this fragmentation, any query about paths connecting two **nodes** can be answered by using just the fragments in which **nodes** are located and the high-speed fragment. In general, if each fragment is managed by a distinguished **processor**, then the query can be answered by three **processors** working in parallel. This schema can be applied recursively to generate an arbitrary number of hierarchical levels. (15 Refs)
  Subfile: C
  Descriptors: database theory; deductive databases; graph theory; parallel algorithms; parallel programming
  Identifiers: **base relation** tuples; graph edges; **base relation partitioning** ; transitive closure queries; parallel computation; semantic data fragmentation; transportation network; fragmentation algorithm; subgraph; high-speed fragment; hierarchical relationship
  Class Codes: C4250  (Database theory); C4240P (Parallel programming and algorithm theory); C1160  (Combinatorial mathematics)

14909393    PASCAL No.: 01-0058742
**Supporting hot spots with** materialized   views
**DaWak 2000 : data warehousing and knowledge discovery** : London, 4-6
**September 2000**

ALBRECHT J; BAUCR A; REDERT M
KAMBAYASHI Yahiko, ed; MOHANIA Mukesh, ed; A MIN TJOA, ed
Department of Database Systems, University of Erlangen-Nuremberg,
Martensstr. 3, 91058 Erlangen, Germany

International conference on data warehousing and knowledge discovery, 2
(London GBR) 2000-09-04

Document Type: P (Serial); C (Conference Proceedings) ; A (Analytic)
Country of Publication: Germany; United States
Language: English

Since  data  warehousing  has  become a major field of research there has
been a lot of interest in the selection of **materialized   views** for query
optimization.  The  problem  is  to  find  the set of **materialized   views**
which  yields  the  highest cost savings for a given set of queries under a
certain  space  constraint.  The  analytical perspective results in queries
which  on the one hand require aggregations but on the other hand are quite
restrictive  with  regard  to the fact data. Usually there are "hot spots",
i.e.  regions  which are requested very frequently, like the current period
or  the  most  important  product  **group** . However,  most  algorithms in
literature  do  not consider restrictions of queries and therefore generate
only  views  containing  all  summary  data  at a certain aggregation level
although  the  space  it  occupies  could  better  be  used for other, more
beneficial  views.  This article presents an algorithm for the selection of
restricted  views.  The  cost  savings  using  this  algorithm  have  been
experimentally  evaluated  to  be up to 80% by supplying only 5% additional
space.

English Descriptors: Query; Database; Multidimensional system; Grain size
   distribution; Aggregate; Grid pattern; **Node** ; Algorithm; Data warehouse

French Descriptors: Question documentaire; Base donnee; Systeme n
   dimensions; Granularite; Agregat; Maillage; Noeud; Algorithme; Vue
   materialisee; Entrepot donnee

```
Set       Items     Description
S1           37     (MATERIALI?)(2N)VIEW?
S2      1067781     GROUP? OR CLUSTER? OR PARTITION? OR CATEGORI?
S3        80338     NODE?
S4            0     S1 AND S2 AND S3
S5            5     S1 AND (S2 OR S3)
```

File 347:JAPIO Nov 1976-2003/Dec(Updated 040402)
        (c) 2004 JPO & JAPIO
File 350:Derwent WPIX 1963-2004/UD,UM &UP=200427
        (c) 2004  Thomson Derwent

DIALOG(R)File 350:Derwent WPIX
(c) 2004  Thomson Derwent. All rts. reserv.

015442269    **Image available**
WPI Acc No: 2003-504411/200347
XRPX Acc No: N03-400560
**Computer-implemented method of maintaining complex** grouping **expression tables of transactions, involves building modification data stream for** materialized view **maintained with table, based on base table modifications**
Patent Assignee: COCHRANE R J (COCH-I); LEHNER W (LEHN-I); PIRAHESH M H (PIRA-I); SIDLE R S (SIDL-I)
Inventor: COCHRANE R J; LEHNER W; PIRAHESH M H; SIDLE R S
Number of Countries: 001  Number of Patents: 001
Patent Family:

| Patent No | Kind | Date | Applicat No | Kind | Date | Week | |
|-----------|------|------|-------------|------|------|------|---|
| US 20030093407 | A1 | 20030515 | US 99135277 | P | 19990521 | 200347 | B |
| | | | US 99453982 | A | 19991202 | | |

Priority Applications (No Type Date): US 99135277 P 19990521; US 99453982 A 19991202
Patent Details:

| Patent No | Kind Lan Pg | Main IPC | Filing Notes |
|-----------|-------------|----------|--------------|
| US 20030093407 A1 | 25 | G06F-007/00 | Provisional application US 99135277 |

Abstract (Basic): US 20030093407 A1
    NOVELTY - A data stream comprising modifications to be propagated to a **materialized view** maintained with complex **grouping** expressions, is built after performing a modification to a base table in a transaction. The built data stream is applied to the **materialized view** .
    DETAILED DESCRIPTION - INDEPENDENT CLAIMS are also included for the following:
    (1) computer program product for maintaining complex **grouping** expression tables of transactions; and
    (2) computer-implemented system for maintaining complex **grouping** expression tables of transactions.
    USE - For maintaining complex **grouping** expression tables of transactions.
    ADVANTAGE - Avoids full re-computation of queries for updating the **materialized views** (i.e. automatic summary tables) of the database tables, by building a data stream comprising table modifications for the **materialized view** and applying the data stream to the view.
    DESCRIPTION OF DRAWING(S) - The figure shows the hardware configuration of the table maintenance system.
    pp; 25 DwgNo 1/9
Title Terms: COMPUTER; IMPLEMENT; METHOD; MAINTAIN; COMPLEX; **GROUP** ; EXPRESS; TABLE; TRANSACTION; BUILD; MODIFIED; DATA; STREAM; VIEW; MAINTAIN; TABLE; BASED; BASE; TABLE; MODIFIED
Derwent Class: T01
International Patent Class (Main): G06F-007/00
File Segment: EPI

DIALOG(R)File 350:Derwent WPIX
(c) 2004  Thomson Derwent. All rts. reserv.

015442269    **Image available**
WPI Acc No: 2003-504411/200347
XRPX Acc No: N03-400560
   **Computer-implemented method of maintaining complex** grouping **expression
   tables of transactions, involves building modification data stream for**
   materialized   view **maintained with table, based on base table
   modifications**
Patent Assignee: COCHRANE R J (COCH-I); LEHNER W (LEHN-I); PIRAHESH M H
   (PIRA-I); SIDLE R S (SIDL-I)
Inventor: COCHRANE R J; LEHNER W; PIRAHESH M H; SIDLE R S
Number of Countries: 001  Number of Patents: 001
Patent Family:

| Patent No | Kind | Date | Applicat No | Kind | Date | Week | |
|-----------|------|------|-------------|------|------|------|---|
| US 20030093407 | A1 | 20030515 | US 99135277 | P | 19990521 | 200347 | B |
| | | | US 99453982 | A | 19991202 | | |

Priority Applications (No Type Date): US 99135277 P 19990521; US 99453982 A
   19991202
Patent Details:

| Patent No | Kind Lan | Pg | Main IPC | Filing Notes |
|-----------|----------|-----|----------|--------------|
| US 20030093407 | A1 | 25 | G06F-007/00 | Provisional application US 99135277 |

Abstract (Basic): US 20030093407 A1
       NOVELTY - A data stream comprising modifications to be propagated
   to a **materialized   view** maintained with complex **grouping**
   expressions, is built after performing a modification to a base table
   in a transaction. The built data stream is applied to the **materialized
   view** .
       DETAILED DESCRIPTION - INDEPENDENT CLAIMS are also included for the
   following:
       (1) computer program product for maintaining complex **grouping**
   expression tables of transactions; and
       (2) computer-implemented system for maintaining complex **grouping**
   expression tables of transactions.
       USE - For maintaining complex **grouping** expression tables of
   transactions.
       ADVANTAGE - Avoids full re-computation of queries for updating the
   **materialized   views** (i.e. automatic summary tables) of the database
   tables, by building a data stream comprising table modifications for
   the **materialized   view** and applying the data stream to the view.
       DESCRIPTION OF DRAWING(S) - The figure shows the hardware
   configuration of the table maintenance system.
       pp; 25 DwgNo 1/9
Title Terms: COMPUTER; IMPLEMENT; METHOD; MAINTAIN; COMPLEX; **GROUP** ;
   EXPRESS; TABLE; TRANSACTION; BUILD; MODIFIED; DATA; STREAM; VIEW;
   MAINTAIN; TABLE; BASED; BASE; TABLE; MODIFIED
Derwent Class: T01
International Patent Class (Main): G06F-007/00
File Segment: EPI

DIALOG(R)File 350:Derwent WPIX
(c) 2004  Thomson Derwent. All rts. reserv.

014928068    **Image available**
WPI Acc No: 2002-748775/200281
   **Method for preparing query using** materialized   view **and dimension
   layer in data warehouse**
Patent Assignee: KOREA ADV INST SCI & TECHNOLOGY (KOAD  )
Inventor: KIM M H; LEE Y J; PARK C S
Number of Countries: 001  Number of Patents: 001
Patent Family:

| Patent No | Kind | Date | Applicat No | Kind | Date | Week |
|-----------|------|------|-------------|------|------|------|

KR 2002046786  A   20020621 KR 200077122   A   20001215 ●200281  B

Priority Applications (No Type Date): KR 200077122 A 20001215
Patent Details:
Patent No  Kind Lan Pg   Main IPC    Filing Notes
KR 2002046786 A      1 G06F-017/30

Abstract (Basic): KR 2002046786 A
        NOVELTY - A method for preparing a query using a **materialized
    view** and a dimension layer in a data warehouse is provided to enhance
    a performance of a data warehouse system by returning the same result
    as a query created by a user and automatically creating new query
    capable of being processed more effectively using many **materialized
    views** for being processed by replacing the original query.
        DETAILED DESCRIPTION - In a preparation method of a query created
    by a user using new query through a dimension layer and many
    **materialized   views** existed in a data warehouse at a data warehouse
    system storing large amount of data, a normalized form is defined with
    respect to the query and the **materialized   views** using a **group**
    grid being induced from the dimension layers(S100). It is checked
    whether each **materialized   view** may be used in the preparation of
    the query(S110). **Materialized   views** are selected to be used in the
    preparation of the query(S120). A query block is created with respect
    to each selected **materialized   view** (S130). New query is created by
    integrating the created query blocks(S170).
        pp; 1 DwgNo 1/10
Title Terms: METHOD; PREPARATION; QUERY; VIEW; DIMENSION; LAYER; DATA;
  WAREHOUSE
Derwent Class: T01
International Patent Class (Main): G06F-017/30
File Segment: EPI


   **5/5/3      (Item 3 from file: 350)**
DIALOG(R)File 350:Derwent WPIX
(c) 2004  Thomson Derwent. All rts. reserv.

014366794     **Image available**
WPI Acc No: 2002-187496/200224
XRPX Acc No: N02-142131
   **Processing method for queries to access** materialized   views **that
   group along an ordered dimension rewriting query if** materialized   view
   **satisfies each condition in set of conditions**
Patent Assignee: ORACLE CORP (ORAC-N)
Inventor: BELLO R; WITKOWSKI A; ZIAUDDIN M
Number of Countries: 001  Number of Patents: 001
Patent Family:
Patent No     Kind   Date    Applicat No   Kind   Date     Week
US 6345272     B1  20020205  US 99361688    A   19990727  200224 B

Priority Applications (No Type Date): US 99361688 A 19990727
Patent Details:
Patent No  Kind Lan Pg   Main IPC    Filing Notes
US 6345272     B1     23 G06F-017/30

Abstract (Basic): US 6345272 B1
        NOVELTY - The method involves receiving an aggregate query that
    places a restriction on an ordered dimension. The restriction is
    specified at a first level of granularity for the dimension. The
    aggregate query does not reference a **materialized   view** that **groups**
     results at a second level of granularity of the ordered dimension. The
    second level of granularity is coarser than the first.
        DETAILED DESCRIPTION - It is determined whether the **materialized
    view** satisfies each condition of a first set of conditions. If the
    **materialized   view** satisfies each condition, then the query is
    rewritten to produce a query that references the **materialized   view**
    and includes a rewritten restriction.
        INDEPENDENT CLAIMS are included for a computer-readable medium and

for a database system.
       USE - For rewriting queries to access a **materialized**   **view** .
       ADVANTAGE - Rewrite mechanism is capable of rewriting queries to
access **materialized**   **views**  that would otherwise not be rewritten.
Does not depend on summary tables.
       DESCRIPTION OF DRAWING(S) - The figure shows a dimension hierarchy.


       pp; 23 DwgNo 1/9
Title Terms: PROCESS; METHOD; QUERY; ACCESS; VIEW; **GROUP** ; ORDER;
   DIMENSION; REWRITING; QUERY; VIEW; SATISFY; CONDITION; SET; CONDITION
Derwent Class: T01
International Patent Class (Main): G06F-017/30
File Segment: EPI



   **5/5/4      (Item 4 from file: 350)**
DIALOG(R)File 350:Derwent WPIX
(c) 2004  Thomson Derwent. All rts. reserv.

014350088     **Image available**
WPI Acc No: 2002-170791/200222
XRPX Acc No: N02-129914
   **Query optimization method for computer system, involves rewriting query**
   **using** materialized   views **having** partitioning **or replication**
   **properties different from properties specified in reference tables of**
   **query**
Patent Assignee: INT BUSINESS MACHINES CORP (IBMC  )
Inventor: COCHRANE R J; LAPIS G; LEUNG T Y; PIRAHESH M H; SIDLE R S; SIMMEN
   D E; URATA M S; ZUZARTE C P
Number of Countries: 001  Number of Patents: 001
Patent Family:
Patent No    Kind   Date     Applicat No   Kind   Date      Week
US 6339769     B1   20020115  US 98152551    A    19980914  200222  B

Priority Applications (No Type Date): US 98152551 A 19980914
Patent Details:
Patent No  Kind Lan Pg   Main IPC    Filing Notes
US 6339769    B1    23 G06F-017/30

Abstract (Basic): US 6339769 B1
       NOVELTY - Existence of multiple **materialized**   **views**  having
   **partitioning**  or replication properties different from the properties
specified in the reference tables in a query, is determined after
accepting the query. Query is rewritten using the **materialized**   **views**
 after analyzing a portion of the query using the **materialized**   **views**
. Rewritten query is executed using the **materialized**   **views** .
       DETAILED DESCRIPTION - INDEPENDENT CLAIMS are also included for the
following:
       (a) Query optimization apparatus;
       (b) a computer program.
       USE - For optimization of queries by transparently altering
properties of relational tables using **materialized**   **views**  in a
database management system.
       ADVANTAGE - Optimizes queries using **materialized**   **views**  that are
replicated and **partitioned**  across multiple processors, and optimizes
RDBMS software using replicated and **partitioned**  copies of
**materialized**   **views** .
       DESCRIPTION OF DRAWING(S) - The figure shows the flowchart
explaining the query optimization method.
       pp; 23 DwgNo 12/12
Title Terms: QUERY; METHOD; COMPUTER; SYSTEM; REWRITING; QUERY; VIEW;
   **PARTITION** ; REPLICA; PROPERTIES; PROPERTIES; SPECIFIED; REFERENCE; TABLE;
   QUERY
Derwent Class: T01
International Patent Class (Main): G06F-017/30
File Segment: EPI

011762733     **Image available**
WPI Acc No: 1998-179643/199816
XRPX Acc No: N98-142080
  **Evaluation method for query having aggregation using** materialised   view
    **- evaluating query using** materialised   view **which is semantically**
  **analysed to determine whether view is usable in evaluating input query**
Patent Assignee: AT & T CORP (AMTT  )
Inventor: DAR S; JAGADISH H V; LEVY A Y; SRIVASTAVA D; JAGADISH H; LEVY A
Number of Countries: 020  Number of Patents: 002
Patent Family:

| Patent No | Kind | Date | Applicat No | Kind | Date | Week | |
|---|---|---|---|---|---|---|---|
| WO 9809238 | A1 | 19980305 | WO 97US14660 | A | 19970819 | 199816 | B |
| US 5897632 | A | 19990427 | US 9624635 | A | 19960827 | 199924 | |
| | | | US 97895024 | A | 19970716 | | |

Priority Applications (No Type Date): US 97895024 A 19970716; US 9624635 P
  19960827
Patent Details:

| Patent No | Kind | Lan | Pg | Main IPC | Filing Notes |
|---|---|---|---|---|---|
| WO 9809238 | A1 | E | 54 | G06F-017/30 | |

    Designated States (National): CA JP
    Designated States (Regional): AT BE CH DE DK ES FI FR GB GR IE IT LU MC
    NL PT SE

| US 5897632 | A | | | G06F-017/30 | Provisional application US 9624635 |
|---|---|---|---|---|---|

Abstract (Basic): WO 9809238 A
      The method involves using **materialised   views** to compute answers
    to SQL queries with **grouping** and aggregation. The **materialised**
    **view** is semantically analysed to determine whether the **materialised**
    **view** can be used in evaluating an input query.
      If the view is usable, the input query is rewritten to produce an
    output query that is multi-set equivalent to the input query, and that
    specifies one or more occurrences of the **materialised   view** as a
    source of information to be returned by the output query. The output
    query is then evaluated.
      USE - Using **materialised   view** to compute and evaluate SQL
    queries with **grouping** and aggregation in query optimisation, data
    warehousing, very large transaction recording systems and mobile
    computing.
      Dwg.1/10
Title Terms: EVALUATE; METHOD; QUERY; AGGREGATE; VIEW; EVALUATE; QUERY;
  VIEW; ANALYSE; DETERMINE; VIEW; EVALUATE; INPUT; QUERY
Derwent Class: T01
International Patent Class (Main): G06F-017/30
File Segment: EPI

```
Set     Items    Description
S1    1067781    CLUSTER? OR GROUP? OR PARTITION? OR CATEGO  ?
S2        167    MATERIAL?()VIEW? OR (STORE? OR SAVE? OR CACHE?)(N)QUER? OR
                 (AUXILIAR? OR BASE)()RELATION?
S3         18    S1 AND S2
S4          7    S3 AND (NODE? OR LOCATION? OR PROCESSOR?)
File 347:JAPIO Nov 1976-2003/Dec(Updated 040402)
         (c) 2004 JPO & JAPIO
File 350:Derwent WPIX 1963-2004/UD,UM &UP=200427
         (c) 2004  Thomson Derwent
```

014928068     **Image available**
WPI Acc No: 2002-748775/200281
  **Method for preparing query using** materialized   view **and dimension layer in data warehouse**
Patent Assignee: KOREA ADV INST SCI & TECHNOLOGY (KOAD  )
Inventor: KIM M H; LEE Y J; PARK C S
Number of Countries: 001  Number of Patents: 001
Patent Family:

| Patent No | Kind | Date | Applicat No | Kind | Date | Week |  |
|-----------|------|------|-------------|------|------|------|--|
| KR 2002046786 | A | 20020621 | KR 200077122 | A | 20001215 | 200281 | B |

Priority Applications (No Type Date): KR 200077122 A 20001215
Patent Details:

| Patent No | Kind | Lan | Pg | Main IPC | Filing Notes |
|-----------|------|-----|----|----------|--------------|
| KR 2002046786 | A | | 1 | G06F-017/30 | |

Abstract (Basic): KR 2002046786 A
      NOVELTY - A method for preparing a query using a **materialized view** and a dimension layer in a data warehouse is provided to enhance a performance of a data warehouse system by returning the same result as a query created by a user and automatically creating new query capable of being processed more effectively using many **materialized views** for being processed by replacing the original query.
      DETAILED DESCRIPTION - In a preparation method of a query created by a user using new query through a dimension layer and many **materialized views** existed in a data warehouse at a data warehouse system storing large amount of data, a normalized form is defined with respect to the query and the **materialized views** using a **group** grid being induced from the dimension layers(S100). It is checked whether each **materialized view** may be used in the preparation of the query(S110). **Materialized views** are selected to be used in the preparation of the query(S120). A query block is created with respect to each selected **materialized view** (S130). New query is created by integrating the created query blocks(S170).
      pp; 1 DwgNo 1/10
Title Terms: METHOD; PREPARATION; QUERY; VIEW; DIMENSION; LAYER; DATA; WAREHOUSE
Derwent Class: T01
International Patent Class (Main): G06F-017/30
File Segment: EP

DIALOG(R)File 350:Derwent ⬤PIX

011762733    **Image available**
WPI Acc No: 1998-179643/199816
XRPX Acc No: N98-142080
  **Evaluation method for query having aggregation using** materialised   view
    **- evaluating query using** materialised   view **which is semantically**
  **analysed to determine whether view is usable in evaluating input query**
Patent Assignee: AT & T CORP (AMTT  )
Inventor: DAR S; JAGADISH H V; LEVY A Y; SRIVASTAVA D; JAGADISH H; LEVY A
Number of Countries: 020  Number of Patents: 002
Patent Family:

| Patent No | Kind | Date | Applicat No | Kind | Date | Week | |
|---|---|---|---|---|---|---|---|
| WO 9809238 | A1 | 19980305 | WO 97US14660 | A | 19970819 | 199816 | B |
| US 5897632 | A | 19990427 | US 9624635 | A | 19960827 | 199924 | |
| | | | US 97895024 | A | 19970716 | | |

Priority Applications (No Type Date): US 97895024 A 19970716; US 9624635 P
  19960827
Patent Details:

| Patent No | Kind | Lan | Pg | Main IPC | Filing Notes |
|---|---|---|---|---|---|
| WO 9809238 | A1 | E | 54 | G06F-017/30 | |

    Designated States (National): CA JP
    Designated States (Regional): AT BE CH DE DK ES FI FR GB GR IE IT LU MC
    NL PT SE

| | | | | | |
|---|---|---|---|---|---|
| US 5897632 | A | | | G06F-017/30 | Provisional application US 9624635 |

Abstract (Basic): WO 9809238 A
      The method involves using **materialised   views** to compute answers
  to SQL queries with **grouping** and aggregation. The **materialised**
  **view** is semantically analysed to determine whether the **materialised**
  **view** can be used in evaluating an input query.
      If the view is usable, the input query is rewritten to produce an
  output query that is multi-set equivalent to the input query, and that
  specifies one or more occurrences of the **materialised   view** as a
  source of information to be returned by the output query. The output
  query is then evaluated.
      USE - Using **materialised   view** to compute and evaluate SQL
  queries with **grouping** and aggregation in query optimisation, data
  warehousing, very large transaction recording systems and mobile
  computing.
      Dwg.1/10
Title Terms: EVALUATE; METHOD; QUERY; AGGREGATE; VIEW; EVALUATE; QUERY;
  VIEW; ANALYSE; DETERMINE; VIEW; EVALUATE; INPUT; QUERY
Derwent Class: T01
International Patent Class (Main): G06F-017/30
File Segment: EPI